arXiv:1606.05689v1 [cs.DS] 17 Jun 2016

Bidimensionality and Kernels*†

Fedor V. Fomin‡ Daniel Lokshtanov† Saket Saurabh†§ Dimitrios M. Thilikos¶‖

Abstract

Bidimensionality theory was introduced by Demaine et al. in 2005 as a tool to obtain *subexponential* time parameterized algorithms on H-minor-free graphs. Demaine and Hajiaghayi extended the theory to obtain polynomial time approximation schemes (PTASs) for bidimensional problems. In this work, we establish a third meta-algorithmic direction for bidimensionality theory by relating it to the existence of linear kernels for parameterized problems. In particular, we prove that every minor (respectively contraction) bidimensional problem that satisfies a separation property and is expressible in Countable Monadic Second Order Logic (CMSO), admits a linear kernel for classes of graphs that exclude a fixed graph (respectively an apex graph) H as a minor. Our results imply that a multitude of bidimensional problems admit linear kernels on the corresponding graph classes. For most of these problems no polynomial kernels on H-minor-free graphs were known prior to our work.

Keywords: Kernelization, Parameterized algorithms, Treewidth, Bidimensionality

1 Introduction

Bidimensionality theory was introduced by Demaine et al. in [18]. This theory is build on cornerstone theorems from Graph Minors Theory of Robertson and Seymour [42] and initially it was developed to unify and extend subexponential fixed-parameter algorithms for NP-hard graph problems to a broad range of graphs including planar graphs, map graphs, bounded-genus graphs and graphs excluding any fixed graph as a minor [16–18, 20, 21]. Roughly speaking, the problem is bidimensional if the solution value for the problem on a $k \times k$-grid is $\Omega(k^2)$, and contraction/removal of edges does not increase solution value. Many natural problems are bidimensional, including Dominating Set, Feedback Vertex Set, Edge Dominating Set, Vertex Cover, r-Dominating Set, Connected Dominating Set, Cycle Packing, Connected Vertex Cover, Graph Metric TSP, and many others.

*Preliminary results of this paper appeared in [30]. Emails of authors: fomin@ii.uib.no, daniello@ii.uib.no, saket@imsc.res.in/Saket.Saurabh@ii.uib.no, sedthilk@thilikos.info.

†The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreements no. 267959 (F. V. Fomin) and no. 306992 (S. Saurabh).

‡Department of Informatics, University of Bergen, Norway.

§Institute of Mathematical Sciences, India.

¶AlGCo project team, CNRS, LIRMM, Montpellier, France.

‖Department of Mathematics, National and Kapodistrian University of Athens, Greece.

The second application of bidimensionality was given by Demaine and Hajiaghayi in [19], where it have been shown that bidimensionality is a useful theory not only in the design of fast fixed-parameter algorithms but also in the design of fast PTASs. Demaine and Hajiaghayi established a link between parameterized and approximation algorithms by proving that every bidimensional problem satisfying some simple separation properties has a PTAS on planar graphs and other classes of sparse graphs. See also [29] for further development of the applications of bidimensionality in the theory of EPTASs We refer to the surveys [21, 23] for further information on bidimensionality and its applications, as well as the book [14].

In this work we give the third application of bidimensionality—*kernelization*. Kernelization can be seen as the strategy of analyzing preprocessing or data reduction heuristics from a parameterized complexity perspective. Parameterized complexity introduced by Downey and Fellows is basically a two-dimensional generalization of "P vs. NP" where in addition to the overall input size n, one studies the effects on computational complexity of a secondary measurement that captures additional relevant information. This additional information can be, for example, solution size or a structural restriction on the input distribution considered, such as the treewidth of an input graph. The secondary information is quantified by a positive integer k and is called the *parameter*. Parameterization can be deployed in many different ways; for general background on the theory see [15, 24, 26, 40].

A parameterized problem with a parameter k is said to admit a *polynomial kernel* if there is a polynomial time algorithm (the degree of polynomial is independent of k), called a *kernelization* algorithm, that reduces the input instance down to an instance with size bounded by a polynomial $p(k)$ in k, while preserving the answer. Kernelization has been extensively studied in parameterized complexity, resulting in polynomial kernels for a variety of problems. Notable examples of known kernels are a $2k$-sized vertex kernel for Vertex Cover [11], a $355k$ kernel for Dominating Set on planar graphs [1], which later was improved to a $67k$ kernel [10], or an $O(k^2)$ kernel for Feedback Vertex Set [43] parameterized by the solution size. One of the most intensively studied directions in kernelization is the study of problems on planar graphs and other classes of sparse graphs. This study was initiated by Alber et al. [2] who gave the first linear-sized kernel for the Dominating Set problem on planar graphs. The work of Alber et al. [2] triggered an explosion of papers on kernelization, and kernels of linear sizes were obtained for a variety of parameterized problems on planar graphs including Connected Vertex Cover, Minimum Edge Dominating Set, Maximum Triangle Packing, Efficient Edge Dominating Set, Induced Matching, Full-Degree Spanning Tree, Feedback Vertex Set, Cycle Packing, and Connected Dominating Set [2, 7, 8, 10, 34, 35, 37, 39]. We refer to the surveys [33, 36, 38] as well as the recent textbook [15] for a detailed treatment of the area of kernelization. Since most of the problems known to have polynomial kernels on planar graphs are bidimensional, the existence of links between bidimensionality and kernelization was conjectured and left as an open problem in [18].

In this work we show that every bidimensional problem with a simple separation property, which is a weaker property than the one required in the framework of Demaine and Hajiaghayi for PTASs [19] and which is expressible in the language of Counting Monadic Second Order Logic (CMSO) (we postpone these definitions till the next section) has a linear kernel on planar and even much more general classes of graphs. In this paper all the problems are parameterized by the

solution size. Our main result is the following theorem.

Theorem 1.1. *Every CMSO-definable linear-separable minor-bidimensional problem* Π *admits a linear kernel on graphs excluding some fixed graph as a minor. Every CMSO-definable linear-separable contraction–bidimensional problem* Π *admits a linear kernel on graphs excluding some fixed apex graph as a minor.*

High level overview of the main proof ideas. Our approach is built on the work of Bodlaender et al. [5] who proved the first meta-theorems on kernelization. In the following we refer to the full version of [5], which is available as [6]. Bodlaender et al. have shown that every parameterized problem that has finite integer index and satisfies an additional surface-dependent property, called a quasi-coverability property, has a linear kernel on graphs of bounded genus.

The kernelization framework in [6] is based on the following idea. Suppose that every yes-instance of a given parameterized problem admits a protrusion decomposition. In other words, suppose that the vertex set of an input graph G can be partitioned in sets $R_0, R_1, \ldots, R_\ell$, where $|R_1|$ and ℓ are of size linear in parameter k, R_0 separates R_i and R_j, $1 \leq i, j \leq \ell$, and every set R_i, $i \geq 1$, is a protrusion, i.e. induces a graph of constant treewidth with a constant number of neighbours in R_0. Then kernelization algorithm uses only one reduction rule, which is based on finite integer index properties of the problem in question, and replaces a protrusion with a protrusion of constant size.

In this paper we use exactly the same reduction rule as the one used in the kernelization algorithm given in [6] for obtaining kernels on planar graphs and graphs of bounded genus. The novel technical contribution of this paper is twofold. First, we introduce the new way the kernel sizes for bidimensional problems are analyzed. The analysis of kernel sizes in [6] requires "topological" decompositions of the given graph, in the sense that the partitioning of the graph into regions with small border, or protrusions, strongly depends on the embedding of a graph into a surface. Then topological properties of embedding are used to prove the existence of a protrusion decomposition. While such an approach works well when we have a topological embedding it seems difficult to extend it to graphs excluding some fixed graph as a minor. Instead of taking the topological approach, we apply bidimensionality and suitable variants of the Excluded Grid Theorem [22, 27]. Roughly speaking, we show that bidimensionality and separability implies existence of a protrusion decomposition. This makes our arguments not only much more general but also considerably simpler than the analysis in [6]. Our second technical contribution is the proof that every CMSO-definiable separable problem has a finite integer index. Pipelined with the framework from [6], these results imply the proof of Theorem 1.1.

The remaining part of the paper is organized as follows. In Section 2, we provide definitions and notations used in the remaining part of the paper. Section 3 is devoted to combinatorial lemmata on separation properties of bidimensional problems. Based on these combinatorial lemmata, we prove a novel decomposition theorem (Theorem 3.11), which is the first main technical contribution of this work. In Section 4, we prove the second main technical contribution of the paper (Theorem 4.4) about the finite integer index of separable CMSO-optimization problems. In Section 5, we prove the main result of the paper about linear kernels. In the concluding Section 6 we discuss the connection between the separable-bidimensional property of a problem and the quasi-coverable property, which

was used in the meta-theorem from [6]. We also provide a brief overview of the relevant results appeared since the conference version of this paper was published.

2 Preliminaries

In this section we give various definitions used in the paper.

Concepts from Graph Theory Let G be a graph. We use the notation $V(G)$ and $E(G)$ for the vertex set and the edge set of G respectively. We say that a graph H is a *subgraph* of G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. Given a set $S \subseteq V(G)$ we denote by $G[S]$ the subgraph G' of G where $V(G') = S$ and $E(G') = \{xy \in E(G) \mid \{x, y\} \subseteq S\}$ and we call G' the *subgraph of G induced by S* or we simply say that G' is an induced subgraph of G.

For every $S \subseteq V$, we denote by $G - S$ the graph obtained from G by removing the vertices of S, i.e. $G - S = G[V(G) \setminus S]$. For vertex $v \in V(G)$ we also use $G - v$ for $G - \{v\}$. For a set $S \subseteq V(G)$, we define $N_G(S)$ to be the *open neighborhood* of S in G, which is the set of vertices from $V(G) \setminus S$ adjacent to vertices of S. The *closed neighborhood* of S is $N_G[S] := N(S) \cup S$. Given a set $S \subseteq V(G)$, we denote by $\partial_G(S)$ the set of all vertices in S that are adjacent in G with vertices not in S. Thus $N_G(S) = \partial_G(V(G) \setminus S)$.

Treewidth. A *tree decomposition* of a graph G is a pair $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$, where T is a tree whose every node t is assigned a vertex subset $X_t \subseteq V(G)$, called a bag, such that the following three conditions hold:

(T1) $\bigcup_{t \in V(T)} X_t = V(G)$. In other words, every vertex of G is in at least one bag.

(T2) For every $uv \in E(G)$, there exists a node t of T such that bag X_t contains both u and v.

(T3) For every $u \in V(G)$, the set $T_u = \{t \in V(T) \ : \ u \in X_t\}$, i.e., the set of nodes whose corresponding bags contain u, induces a connected subtree of T.

The *width* of a tree decomposition $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$ equals $\max_{t \in V(T)} |X_t| - 1$, that is, the maximum bag size, minus one. The *treewidth* of a graph G, denoted by $\mathrm{tw}(G)$, is the minimum possible width of a tree decomposition of G. To distinguish between the vertices of the decomposition tree T and the vertices of the graph G, we will refer to the vertices of T as *nodes*.

Separators and separations Let G be a graph, $Q \subseteq V(G)$, and let $A_1, A_2 \subseteq V(G)$ such that $A_1 \cup A_2 = V(G)$. We say that the pair (A_1, A_2) is a *separation* of G if there is no edge with one endpoint in $A_1 \setminus A_2$ and the other in $A_2 \setminus A_1$. The *order* of a separation (A_1, A_2) is $|A_1 \cap A_2|$. For a vertex subset $Q \subseteq V(G)$ we say that a separation (A_1, A_2) is a *2/3-balanced separation* of (G, Q) if each of the parts $A_1 \setminus A_2$ and $A_2 \setminus A_1$ contains at most $\frac{2}{3}|Q|$ vertices of Q.

The following separation property of graphs of small treewidth is well known, see e.g. [4, 14].

Proposition 2.1. *Let G be a graph and let $S \subseteq V(G)$. There is a 2/3-balanced separation (A_1, A_2) of (G, S) of order at most $\mathrm{tw}(G) + 1$.*

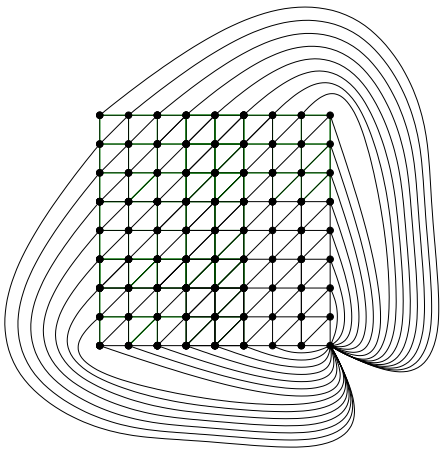

Figure 1: Graph Γ_9.

The following is an easy fact about treewidth, see e.g. [4].

Proposition 2.2. *The treewidth of a graph is the maximum treewidth of its connected components.*

Minors and contractions Given an edge $e = xy$ of a graph G, the graph G/e is obtained from G by contracting the edge e, that is, the endpoints x and y are replaced by a new vertex $v_{x,y}$ which is adjacent to the old neighbors of x and y (except from x and y). A graph H obtained by a sequence of edge-contractions is said to be a *contraction* of G. We denote it by $H \leq_c G$. A graph H is a *minor* of a graph G if H is the contraction of some subgraph of G and we denote it by $H \leq_m G$. We say that a graph G is *H-minor-free* when it does not contain H as a minor. We also say that a graph class $\mathcal{G}$ is *H-minor-free* (or, excludes H as a minor) when all its members are H-minor-free. A graph G is an *apex graph* if there exists a vertex v such that $G - v$ is planar. A graph class $\mathcal{G}$ is *apex-minor-free* if there exists an apex graph H that is not in $\mathcal{G}$.

A graph class $\mathcal{G}$ is said to be *subgraph-closed* (*minor-closed/contraction-closed*) if every subgraph (minor/contraction) of a graph in $\mathcal{G}$ also belongs to $\mathcal{G}$.

Grids and triangulated grids. Given a positive integer k, we denote by $\boxplus_k$ the $(k \times k)$-grid. Formally, for a positive integer k, a *$k \times k$ grid* $\boxplus_k$ is a graph with vertex set $\{(x, y) \ : \ x, y \in \{1, \ldots, k\}\}$. Thus $\boxplus_k$ has exactly k^2 vertices. Two different vertices (x, y) and (x', y') are adjacent if and only if $|x - x'| + |y - y'| = 1$.

For an integer $t > 0$, the graph Γ_t is obtained from the grid $\boxplus_t$ by adding, for all $1 \leq x, y \leq t-1$, the edge $(x+1, y), (x, y+1)$, and additionally making vertex (t, t) adjacent to all the other vertices (x, y) with $x \in \{1, t\}$ or $y \in \{1, t\}$, i.e., to the whole border of $\boxplus_t$. Graph Γ_9 is shown in Fig. 1.

We also need the following result of Robertson and Seymour [41].

Proposition 2.3. *For every $t \geq 1$, $\mathrm{tw}(\boxplus_t) = t$.*

Parameterized graph problems. A parameterized graph problem Π in general can be seen as a subset of $\Sigma^* \times \mathbb{Z}^+$ where, in each instance (x, k) of Π, x encodes a graph and k is the parameter

(we denote by $\mathbb{Z}^+$ the set of all non-negative integers). In this paper we use an extension of this definition used in [6] that permits the parameter k to be negative with the additional constraint that either all pairs with non-positive value of the parameter are in Π or that no such pair is in Π. Formally, a parametrized problem Π is a subset of $\Sigma^* \times \mathbb{Z}$ where for all $(x_1, k_1), (x_2, k_2) \in \Sigma^* \times \mathbb{Z}$ with $k_1, k_2 < 0$ it holds that $(x_1, k_1) \in \Pi$ if and only if $(x_2, k_2) \in \Pi$. This extended definition encompasses the traditional one and is being adopted for technical reasons (see Subsection 4.3). In an instance of a parameterized problem (x, k), the integer k is called the parameter.

The notion of *kernelization* is formally defined as follows.

Definition 2.4. A *kernelization algorithm*, or simply a *kernel*, for a parameterized problem Π is an algorithm $\mathcal{A}$ that, given an instance (x, k) of Π, works in polynomial-time and returns an equivalent instance (x', k') of Π. Moreover, there exists a computable function $g(\cdot)$ such that whenever (x', k') is the output for an instance (x, k), then it holds that $|x'| + k' \leq g(k)$. If the upper bound $g(\cdot)$ is a polynomial (linear) function of the parameter, then we say that Π admits a *polynomial (linear) kernel*.

We often abuse the notation and call the output of a kernelization algorithm, the "reduced" equivalent instance, also a kernel.

Bidimensionality We will restrict our attention to vertex or edge subset problems. A *vertex subset problem* Π is a parameterized problem where input is a graph G and an integer k, the parameter is k. An instance (G, k) is a yes-instance if and only if there exists a set $S \subseteq V(G)$ such that $|S| \leq k$ for minimization prolems (or $|S| \geq k$ for maximization problems) so that a predicate $\phi(G, S)$ is true. Here ϕ can be any computable function which takes as input a graph G and set $S \subseteq V(G)$ and outputs **true** or **false**.

The interpretation is that ϕ defines the space of *feasible solutions* S for a graph G by returning whether S is feasible for G. For an example, for the Dominating Set problem we have that $\phi(G, S) =$ **true** if and only if $N[S] = V(G)$. *Edge subset problems* are defined similarly, with the only difference being that S is a subset of $E(G)$ rather than $V(G)$.

Let us remark that there are many vertex/edge subset problems which at a first glance do not look as if they could be captured by this definition. An example is the Cycle Packing problem. Here input is a graph G and integer k, and the task is to determine whether G contains k pairwise vertex-disjoint cycles $C_1, C_2, \ldots C_k$. This is a vertex subset problem because G has k vertex-disjoint cycles if and only if there exists a set $S \subseteq V(G)$ of size at least k and $\phi(G, S)$ is true, where $\phi(G, S)$ is defined as follows.

$\phi(G, S) = \exists$ subgraph G' of G such that
- each connected component of G' is a cycle,
- and each connected component of G' contains exactly one vertex of S.

This definition may seem a bit silly, since checking whether $\phi(G, S)$ is true for a given graph G and set S is NP-complete. In fact this problem is considered as a more difficult problem than

Cycle Packing. Nevertheless, this definition shows that Cycle Packing is a vertex subset problem, which will allow us to give a linear kernel for Cycle Packing on minor free graphs.

For any vertex or edge subset minimization problem Π we have that $(G,k) \in \Pi$ implies that $(G,k') \in \Pi$ for all $k' \geq k$. Similarly, for a vertex or edge subset maximization problem we have that $(G,k) \in \Pi$ implies that $(G,k') \in \Pi$ for all $k' \leq k$. Thus the notion of "optimality" is well defined for vertex and edge subset problems.

Definition 2.5. For a vertex/edge subset minimization problem Π, we define

$$OPT_\Pi(G) = \min\{k \ : \ (G,k) \in \Pi\}.$$

If no k such that $(G,k) \in \Pi$ exists, $OPT_\Pi(G)$ returns $+\infty$. For a vertex/edge subset maximization problem Π,

$$OPT_\Pi(G) = \max\{k \ : \ (G,k) \in \Pi\}.$$

If no k such that $(G,k) \in \Pi$ exists, $OPT_\Pi(G)$ returns $-\infty$. We define $SOL_\Pi(G)$ to be a function that given as an input a graph G returns a set S of size $OPT_\Pi(G)$ such that $\phi(G,S) = \textbf{true}$, and returns **null** if no such set S exists.

For many problems it holds that contracting an edge can not increase the size of the optimal solution. We will say that such problems are contraction closed. Formally we have the following definition.

Definition 2.6. A vertex/edge subset problem Π is *contraction-closed* if for any G and $uv \in E(G)$, $OPT_\Pi(G/uv) \leq OPT_\Pi(G)$.

If contracting edges, deleting edges and deleting vertices can not increase the size of the optimal solution, we say that the problem is minor-closed.

Definition 2.7. A vertex/edge subset problem Π is *minor-closed* if for any G, edge $uv \in E(G)$ and vertex $w \in V(G)$, $OPT_\Pi(G/uv) \leq OPT_\Pi(G)$, $OPT_\Pi(G \setminus uv) \leq OPT_\Pi(G)$, and $OPT_\Pi(G - w) \leq OPT_\Pi(G)$.

We are now ready to give the definition of bidimensional problems.

Definition 2.8 (Bidimensional problem). A vertex/edge subset problem Π is

- *contraction-bidimensional* if it is contraction-closed, and there exists a constant $c > 0$ such that $OPT_\Pi(\Gamma_k) \geq ck^2$.
- *minor-bidimensional* if it is minor-closed, and there exists a constant $c > 0$ such that $OPT_\Pi(\boxplus_k) \geq ck^2$.

It is usually quite easy to determine whether a problem is contraction (or minor)-bidimensional. Take for an example Independent Set. Contracting an edge may never increase the size of the maximum independent set, so the problem is contraction closed. Furthermore, in Γ_k, the vertex set

$$\{(x,y) \ : \ x \equiv k+1 \mod 2 \ \ \textit{and} \ \ y \equiv k+1 \mod 2\}$$

forms an independent set of size $\frac{(k-1)^2}{4}$. Thus Independent Set is contraction-bidimensional. On the other hand deleting edges may increase the size of a maximum size independent set in G. Thus Independent Set is not minor-bidimensional.

Counting Monadic Second Order Logic The syntax of Monadic Second Order Logic (MSO) of graphs includes the logical connectives $\vee$, $\wedge$, $\neg$, $\Leftrightarrow$, $\Rightarrow$, variables for vertices, edges, sets of vertices, and sets of edges, the quantifiers $\forall$, $\exists$ that can be applied to these variables, and the following five binary relations:

1. $u \in U$ where u is a vertex variable and U is a vertex set variable;
2. $d \in D$ where d is an edge variable and D is an edge set variable;
3. $\mathbf{inc}(d, u)$, where d is an edge variable, u is a vertex variable, and the interpretation is that the edge d is incident with the vertex u;
4. $\mathbf{adj}(u, v)$, where u and v are vertex variables and the interpretation is that u and v are adjacent;
5. equality of variables representing vertices, edges, sets of vertices, and sets of edges.

In addition to the usual features of monadic second-order logic, if we have atomic sentences testing whether the cardinality of a set is equal to q modulo r, where q and r are integers such that $0 \leq q < r$ and $r \geq 2$, then this extension of the MSO is called *counting monadic second-order logic*. Thus CMSO is MSO enriched with the following atomic sentence for a set S:

$\mathbf{card}_{q,r}(S) = \mathbf{true}$ if and only if $|S| \equiv q \pmod{r}$.

We refer to [3, 12, 13] for a detailed introduction on CMSO.

We consider CMSO sentences evaluated either on graphs or on annotated graphs. A class of graphs $\mathcal{G}$ is *CMSO-definable* if there is a CMSO sentence on graphs ϕ such that $G \in \mathcal{G}$ if and only if $G \models \phi$. Similarly a predicate ϕ on annotated graphs is *CMSO-definable* if there is a CMSO sentence on annotated graphs ψ such that $\phi(G, S) = \mathbf{true}$ if and only if $(G, S) \models \phi$.

A vertex/edge subset minimization (or maximization) problem with feasibility function ϕ is a MIN-CMSO problem (or MAX-CMSO problem) if there exists a CMSO sentence ψ such that $\phi(G, S) = \mathbf{true}$ if and only if $(G, S) \models \psi$.

Subquadratic grid minor property In general it is known that there exists a constant c such that any graph G which excludes a $\boxplus_t$ as a minor has treewidth at most $O(t^c)$. I. The exact value of c remains unknown, but it is at least 2 and at most 99 [9]. We will restrict our attention to graph classes on which $c < 2$. In particular we say that a graph class $\mathcal{G}$ has the *subquadratic grid minor property* (SQGM property for short) if there exist constants $\lambda > 0$ and $1 \leq c < 2$ such that any graph $G \in \mathcal{G}$ which excludes $\boxplus_t$ as a minor has treewidth at most λt^c.

Problems that are contraction-closed but not minor closed are considered on more restricted classes of graphs. We say that a graph class $\mathcal{G}$ has the *subquadratic gamma contraction* (SQGC property for short) if there exist constants $\lambda > 0$ and $1 \leq c < 2$ such that any *connected* graph $G \in \mathcal{G}$ which excludes a Γ_t as a contraction has treewidth at most λt^c. Since Γ_t contains $\boxplus_t$ as a minor (in fact as a subgraph), we obtain the following observation.

Observation 2.9. *Every graph class G with the SQGC property has the SQGM property.*

The following proposition follows directly from the linearity of excluded grid-minor in H-minor-free graphs proven by Demaine and Hajiaghayi [22] and its analog for contraction-minors from [27].

Proposition 2.10. *For every graph H, H-minor-free graph class $\mathcal{G}$ has the SQGM property with $c = 1$. If H is an apex graph, then $\mathcal{G}$ has the SQGC property with $c = 1$.*

We say that a parameterized problem Π is a problem on the graph class $\mathcal{G}$ if every yes-instance (G, k) of Π satisfies $G \in \mathcal{G}$. The *restriction* of a parameterized problem Π to a graph class $\mathcal{G}$, denoted by $\Pi \Cap \mathcal{G}$ is defined as follows.

$$\Pi \Cap \mathcal{G} = \{(G, k) \mid (G, k) \in \Pi \text{ and } G \in \mathcal{G}\}.$$

For a parameterized problem Π (on general graphs) we will refer to the restriction of Π to $\mathcal{G}$ by "Π on $\mathcal{G}$".

3 Decomposing into protrusions

In this section we give the main technical combinatorial contribution of this work establishing a protrusion decomposition theorem for linealrly-separable bidimensional problems. The proof of the theorem is done in two steps. First we show that every graph class with SQGM or SQDC property admit a treewidth-modulator of size linear in the parameter of a linealrly-separable (minor- or contraction-) bidimensional problem (Subsection 3.2). Then in Subsection 3.3, we show that graph classes with SQGM and SQDC properties and having linear treewidth-modulators, can be decomposed into protrusions.

3.1 Parameter-treewidth bounds

The following lemmata establish *parameter-treewidth bounds*, that is tight relationships between the size of the optimal solution and the treewidth of the input graph. This relationship was first observed by Demaine et al. [18]. The bound for contraction-bidimensional problems presented here is essentially identical to the one presented in Fomin et al. [27]. We re-prove the lemmata here because of slight differences in definitions.

Lemma 3.1. *For any minor-bidimensional problem Π on a graph class G with the SQGM property, there exist constants $0 < \alpha$, $0 < \epsilon < 1$ such that for any graph $G \in \mathcal{G}$, $\mathrm{tw}(G) \leq \alpha \cdot (OPT_\Pi(G))^\epsilon$.*

Proof. Let λ and c be the constants from the definition of the SQGM property, that is any graph $G \in \mathcal{G}$ which excludes a $\boxplus_t$ as a minor has treewidth at most λt^c. Let β be the constant from the definition of minor-bidimensionality of Π, i.e $OPT_\Pi(\boxplus_t) \geq \beta t^2$. Consider now a graph $G \in \mathcal{G}$. Let t be the maximum integer such that G contains $\boxplus_t$ as a minor. We have that $\mathrm{tw}(G) < \lambda(t+1)^c$. Rearranging terms yields that $(\frac{\mathrm{tw}(G)}{\lambda})^{1/c} < t + 1$ implying that $t \geq (\frac{\mathrm{tw}(G)}{\lambda})^{1/c}$. Since Π is minor-closed it follows that $OPT_\Pi(G) \geq OPT_\Pi(\boxplus_t)$, and since Π is minor-bidimensional we have that

$$OPT_\Pi(G) \geq OPT_\Pi(\boxplus_t) \geq \beta t^2 \geq \beta(\frac{\mathrm{tw}(G)}{\lambda})^{2/c}.$$

Hence

$$\mathrm{tw}(G) \leq \frac{\lambda}{\beta^{\frac{c}{2}}} \cdot OPT_{\Pi}(G)^{\frac{c}{2}}.$$

Since $c < 2$ in the definition of the SQGM property, the statement of the lemma follows. □

Lemma 3.2. *For any contraction-bidimensional problem* Π *on a graph class* G *with the SQGC property, there exist constants* $0 < \alpha$, $0 < \epsilon < 1$, *such that for any connected graph* $G \in \mathcal{G}$, $\mathrm{tw}(G) \leq \alpha \cdot (OPT_{\Pi}(G))^{\epsilon}$.

The proof of the Lemma 3.2 is identical to the proof of Lemma 3.1. The differences between Lemmata 3.1 and 3.2 are as follows. Lemma 3.1 is for minor-closed problems, on graph classes $\mathcal{G}$ with the SQGM property and works for *every* $G \in \mathcal{G}$. Lemma 3.2 is for contraction-closed problems, on graph classes $\mathcal{G}$ with the SQGC property and works for *connected* graphs $G \in \mathcal{G}$.

3.2 Separability and treewidth modulators

We now restrict our attention to problems Π that are somewhat well-behaved in the sense that whenever we have a small separator in the graph that splits the graph in two parts L and R, the intersection $|X \cap L|$ of L with any optimal solution X to the entire graph is a good estimate of $OPT_{\Pi}(G[L])$. This restriction allows us to prove decomposition theorems which are very useful for giving kernels. Similar decomposition theorems may also be used to give approximation schemes, see [19, 29].

Definition 3.3 (Separability). Let $f : \mathbb{Z}^+ \to \mathbb{Z}^+$ be a function. We say that a vertex/edge subset problem Π is *f-separable* if for any graph G and subset $L \subseteq V(G)$ such that $|\partial_G(L)| \leq t$, it holds that

$$|SOL_{\Pi}(G) \cap L| - f(t) \leq OPT_{\Pi}(G[L]) \leq |SOL_{\Pi}(G) \cap L| + f(t).$$

Π is called *separable* if there exists a function f such that Π is *f-separable*. Π is called *linear-separable* if there exists a contant c such that Π is $c \cdot t$-separable.

A nice feature of separable contraction-bidimensional problems is that it is possible to extend Lemma 3.1 to disconnected graphs, since the issue of different connected components influencing the value of the optimal solution disappears.

Lemma 3.4. *For any contraction-bidimensional separable problem* Π *on a graph class* $\mathcal{G}$ *with the SQGC property, there exist constants* $0 < \alpha$, $0 < \epsilon < 1$ *such that for any graph* $G \in \mathcal{G}$, $\mathrm{tw}(G) \leq \alpha \cdot (OPT_{\Pi}(G))^{\epsilon}$.

Proof. Since Π is separable there exists a constant c such that for any graph G and connected component C of G, it holds that

$$OPT_{\Pi}(G[C]) \leq |SOL_{\Pi}(G) \cap C| + c \leq OPT_{\Pi}(G) + c.$$

By Lemma 3.1, there exist constants α and $\epsilon < 1$ such that $\mathrm{tw}(G[C]) \leq \alpha \cdot OPT_{\Pi}(G[C])^{\epsilon}$ for every component C. Since the treewidth of G is at most the maximum of the treewidth of its connected components, the lemma follows by increasing α by a factor c. □

We will consider a sequence of "canonical" bidimensional problems, with one problem for every $\eta \geq 0$. We say that a set $S \subseteq V(G)$ is a *treewidth-η-modulator* if $\mathrm{tw}(G - S) \leq \eta$. We define the following problem.

TREEWIDTH-η-MODULATOR	
Instance:	A graph G, and integer $k \geq 0$.
Parameter:	k.
Problem:	Decide whether there exists a treewidth-η-modulator of G of size k.

It is easy to see that the TREEWIDTH-η-MODULATOR problem is minor-closed. By Proposition 2.3, every $(\eta + 1) \times (\eta + 1)$ subgrid of $\boxplus_t$ must contain at least one vertex of any solution, therefore

$$OPT_\Pi(\boxplus_t) \geq \lfloor (\frac{t}{\eta + 1}) \rfloor^2.$$

Thus the problem is minor-bidimensional.

To see that the problem is separable, consider any graph G, and vertex subset $L \subseteq V(G)$ with $\partial(L) = t$. For any treewidth-η-modulator S of G we have that $(S \cap L) \cup \partial(L)$ is a treewidth-η-modulator of $G[L]$. This proves $OPT_\Pi(G[L]) \leq |SOL_\Pi(G) \cap L| + t$. On the other hand, for any optimal treewidth-η-modulator S of G and treewidth-η-modulator S_L of $G[L]$, we have that $|S \cap L| \leq |S_L| + t$, since otherwise $(S \setminus L) \cup S_L \cup \partial(L)$ is a treewidth-η-modulator of G of size strictly smaller than $|S|$, contradicting the optimality of S. It follows that $|SOL_\Pi(G) \cap L| - t \leq OPT_\Pi(G[L])$. This shows that the TREEWIDTH-η-MODULATOR problem is minor-bidimensional and separable.

The fact that the TREEWIDTH-η-MODULATOR problem is minor-bidimensional and separable, together with Lemma 3.1 yields the following observation.

Observation 3.5. *For every graph class $\mathcal{G}$ with the SQGM property and every $\eta \geq 0$ there exist constants $\alpha \geq 0$ and $0 \leq \lambda < 1$ such that any graph $G \in \mathcal{G}$ with a treewidth-η-modulator S has treewidth at most $\alpha |S|^\lambda$.*

Next we show that the TREEWIDTH-η-MODULATOR problems are canonical bidimensional problems in the following sense.

Lemma 3.6. *For any real $\epsilon > 0$ and minor-bidimensional linear-separable problem Π on graph class $\mathcal{G}$ with the SQGM property, there exists an integer $\eta \geq 0$ such that any graph $G \in \mathcal{G}$ has a treewidth-η-modulator S of size at most $\epsilon \cdot OPT_\Pi(G)$.*

Proof. Let β be a constant such that Π is $(\beta \cdot t) - separable$. Let α' and $0 \leq \lambda < 1$ be the constants from Lemma 3.1, in particular $\mathrm{tw}(G) \leq \alpha' \cdot (OPT_\Pi(G))^\lambda$. Set $\alpha = \max\{\alpha', 1\}$. Furthermore, if $\beta < 1$ then Π is t-separable, and so we may assume without loss of generality that $\beta \geq 1$.

We now define a few constants. The reason these constants are defined the way they are will become clear during the course of the proof. Finally we set η based on α, β, λ and ϵ.

- Set $\rho = \frac{1^\lambda + 2^\lambda - 3^\lambda}{3^\lambda}$ and note that $\rho > 0$.
- set $\gamma = 4\alpha\beta$,
- set $\delta = \frac{\gamma(2\epsilon+1)}{\rho}$,

- set $k_0 = (3+3\gamma)^{\frac{1}{1-\lambda}} + (\frac{\delta}{\epsilon})^{\frac{1}{1-\lambda}} \cdot \frac{1}{3}$. It is easy to verify that k_0 satisfies $\frac{2}{3}k_0 + \gamma k_0^\lambda \leq k_0 - 1$ and $0 \leq \epsilon\frac{1}{3}k_0 - \delta\left(\frac{1}{3}k_0\right)^\lambda$. In fact the first term ensures that the first inequality holds, while the second term ensures the correctness of the second inequality. Finally,
- set $\eta = \alpha \cdot k_0^\lambda$.

We prove by induction on k, that for any $k \geq \frac{1}{3}k_0$, any graph $G \in \mathcal{G}$ such that $OPT_\Pi(G) \leq k$ has a treewidth-η-modulator of size at most $\epsilon k - \delta k^\lambda$. In the base case we consider any k such that $\frac{1}{3}k_0 \leq k \leq k_0$. By Lemma 3.1, any graph $G \in \mathcal{G}$ such that $OPT_\Pi(G) \leq k_0$ has treewidth at most $\alpha \cdot k_0^\lambda = \eta$. Thus G has a treewidth-η-modulator of size 0, and

$$0 \leq \epsilon\frac{1}{3}k_0 - \delta\left(\frac{1}{3}k_0\right)^\lambda \leq \epsilon k - \delta k^\lambda$$

by the choice of k_0. In the last inequality we used that for any $0 \leq \lambda < 1$, ϵ and δ the function $\epsilon k - \delta k^\lambda$ is monotonically increasing from the first point where it becomes positive. This fact may easily be verified by differentiation.

For the inductive step, let $k > k_0$ and suppose that the statement is true for all values below k. We prove the statement for k. Consider a graph $G \in \mathcal{G}$ such that $OPT_\Pi(G) \leq k$. By Lemma 3.1, the treewidth of G is at most $\mathrm{tw}(G) \leq \alpha \cdot k^\lambda$. By Lemma 2.1 applied to $(G, SOL_\Pi(G))$ there is a 2/3-balanced separation (A_1, A_2) of $(G, SOL_\Pi(G))$ of order at most $\mathrm{tw}(G)+1 \leq \alpha \cdot k^\lambda + 1$. Let $L = A_1 \setminus A_2$, $S = A_1 \cap A_2$ and $R = A_2 \setminus A_1$. Note that there are no edges from L to R. Since (A_1, A_2) is a 2/3-balanced separation it follows that there exists a real $\frac{1}{3} \leq a \leq \frac{2}{3}$ such that $|L \cap SOL_\Pi(G)| \leq a|SOL_\Pi(G)|$ and $|R \cap SOL_\Pi(G)| \leq (1-a)|SOL_\Pi(G)|$.

Consider now the graph $G[L \cup S]$. Since L has no neighbors in R (in G) and Π is $(\beta \cdot t)$-separable, it follows that

$$\begin{aligned} OPT_\Pi(G[L \cup S]) &\leq |SOL_\Pi(G) \cap (L \cup S)| + \beta|S| \\ &\leq ak + (\alpha k^\lambda + 1) + \beta(\alpha k^\lambda + 1) \\ &\leq ak + (\alpha k^\lambda + 1)(\beta + 1) \leq ak + \gamma k^\lambda. \end{aligned}$$

Here the last inequality follows from the assumption that $k \geq k_0 \geq 1$ and the choice of γ. Since $k > k_0$, the properties of k_0 imply that $\frac{2}{3}k + \gamma k^\lambda \leq k - 1$. Further $ak + \gamma k^\lambda \geq \frac{1}{3}k_0$ since $a \geq \frac{1}{3}$. Thus we may apply the induction hypothesis to $G[L \cup S]$ and obtain a treewidth-η-modulator Z_L of $G[L \cup S]$, such that

$$\begin{aligned} |Z_L| &\leq \epsilon(ak + \gamma k^\lambda) - \delta\left(ak + \gamma k^\lambda\right)^\lambda \\ &\leq \epsilon(ak + \gamma k^\lambda) - \delta k^\lambda a^\lambda. \end{aligned}$$

An identical argument applied to $G[R \cup S]$ yields a treewidth-η-modulator Z_R of $G[R \cup S]$, such that

$$|Z_R| \leq \epsilon\left((1-a)k + \gamma k^\lambda\right) - \delta k^\lambda (1-a)^\lambda.$$

We now make a treewidth-η-modulator Z of G as follows. Let $Z = Z_L \cup S \cup Z_R$. The set Z is a treewidth-η-modulator of G because every connected component of $G - Z$ is a subset of L or

R, and Z_L and Z_R are treewidth-η-modulators for $G[L \cup S]$ and $G[R \cup S]$ respectively. Finally we bound the size of Z.

$$\begin{aligned}
|Z| &\leq |Z_L| + |Z_R| + |S| \\
&\leq \epsilon(ak + \gamma k^\lambda) - \delta k^\lambda a^\lambda + \epsilon\left((1-a)k + \gamma k^\lambda\right) - \delta k^\lambda (1-a)^\lambda + \gamma k^\lambda \\
&= \epsilon k - \delta k^\lambda \left((1-a)^\lambda + a^\lambda\right) + k^\lambda \gamma(2\epsilon + 1) \\
&\leq \epsilon k - \delta k^\lambda + k^\lambda \left(\gamma(2\epsilon+1) - \delta\rho\right) \\
&\leq \epsilon k - \delta k^\lambda
\end{aligned}$$

In the transition from the third to the fourth line we used that $(1-a)^\lambda + a^\lambda - 1 \geq \rho$ for any a between $\frac{1}{3}$ and $\frac{2}{3}$.

To conclude the proof, we observe that the statement of the lemma follows from what has just been proved. If $OPT_\Pi(G) \leq k_0$ then G has a treewidth-η-modulator of size $\epsilon \cdot OPT_\Pi(G)$. If $OPT_\Pi(G) > k_0$ then G has a treewidth-η-modulator of size at most $\epsilon \cdot OPT_\Pi(G) - \delta(OPT_\Pi(G))^\lambda$. This completes the proof. □

An identical argument yields an analogous lemma for contraction-bidimensional problems.

Lemma 3.7. *For any real $\epsilon > 0$ and contraction-bidimensional linear-separable problem Π on graph class $\mathcal{G}$ with the SQGC property, there exists an integer $\eta \geq 0$ such that any graph $G \in \mathcal{G}$ has a treewidth-η-modulator S of size at most $\epsilon \cdot OPT_\Pi(G)$.*

3.3 Protrusion decomposition

The notions of *protrusion* and *protrusion decomposition* were introduce in [5], see also [6].

Definition 3.8. [**t-protrusion**] For a graph G, a set $X \subseteq V(G)$ is a *t-protrusion* of G if $|\partial(X)| \leq t$ and $\mathrm{tw}(G[X]) \leq t$.

Definition 3.9. [**(α, r)-Protrusion decomposition**] An *(α, r)-protrusion decomposition* of a graph G is a partition $\mathcal{P} = \{R_0, R_1, \ldots, R_\ell\}$ of $V(G)$ such that

- $\max\{\ell, |R_0|\} \leq \alpha$,
- each $R_i^+ = N_G[R_i]$, $i \in \{1, \ldots, \ell\}$, is a r-protrusion of G, and
- for every $i \in \{1, \ldots, \ell\}$, $N_G(R_i) \subseteq R_0$.

We call the sets R_i^+, $i \in \{1, \ldots, \ell\}$, the *protrusions* of $\mathcal{P}$ and the set R_0 the *core* of $\mathcal{P}$.

Next we prove that the existence of a treewidth-η-modulator S implies the existence of a protrusion decomposition of G into r-protrusions with a core not much bigger than S, and r only depending on η (and the graph class $\mathcal{G}$ we are working with). The inductive proof of the following lemma is similar to the proof of Lemma 3.6 about treewidth-η-modulator, however this time we use induction to construct the required protrusion decomposition.

Lemma 3.10. *Let $\mathcal{G}$ be a graph class with the SQGM property. For any real $\epsilon > 0$ and positive integer η, there exists an integer r such that if $G \in \mathcal{G}$ has a treewidth-η-modulator $S \neq \emptyset$, then G has $((1+\epsilon)|S|, r)$-protrusion decomposition $\mathcal{P}$ with S contained in the core of $\mathcal{P}$.*

Proof. Let α' and $0 \leq \lambda < 1$ be the constants from Observation 3.5, in particular any graph in $\mathcal{G}$ with a treewidth-η-modulator of size at most k has treewidth at most $\alpha' \cdot k^\lambda$. Set $\alpha = \max\{\alpha', 1\}$. We set r based on η, α, λ and ϵ, which are defined as follows.

- Set $\rho = \frac{1^\lambda + 2^\lambda - 3^\lambda}{3^\lambda}$ and note that $\rho > 0$.
- set $\delta = \frac{(1+\epsilon)4\alpha}{\rho}$,
- set $k_0 = (3+6\alpha)^{\frac{1}{1-\lambda}} + (\frac{\delta}{\epsilon})^{\frac{1}{1-\lambda}} \cdot \frac{1}{3}$. Observe that k_0 satisfies $\frac{2}{3}k_0 + 2\alpha k_0^\lambda \leq k_0 - 1$ and $0 \leq \epsilon\frac{1}{3}k_0 - \delta\left(\frac{1}{3}k_0\right)^\lambda$. The first term ensures that the first inequality holds, while the second term ensures the correctness of the second inequality. Finally,
- set $r = \max\{k_0, \alpha k_0^\lambda, \eta\}$

We prove by induction on k the following. For every $k \geq \frac{1}{3}k_0$ if a graph $G \in \mathcal{G}$ has a treewidth-η-modulator S of size at most k, then G has $((1+\epsilon)k - \delta k^\lambda, r)$-protrusion decomposition. In the base case we consider any k such that $\frac{1}{3}k_0 \leq k \leq k_0$. Any graph that has a treewidth-η-modulator S of size at most k_0 has treewidth at most $\alpha \cdot k_0^\lambda \leq r$ by Observation 3.5. Make a protrusion-decomposition with core S and only one protrusion, namely $V(G) \setminus S$. Since $k_0 \leq r$ and $k \geq 1$, this is a (k, r)-protrusion decomposition. To complete the proof of the base case we need to show that $k \leq (1+\epsilon)k - \delta k^\lambda$, or equivalently that $0 \leq \epsilon k - \delta k^\lambda$. However, we have that

$$0 \leq \epsilon\frac{1}{3}k_0 - \delta\left(\frac{1}{3}k_0\right)^\lambda \leq \epsilon k - \delta k^\lambda$$

by the choice of k_0. In the last inequality we used that for any $0 \leq \lambda < 1$, $\epsilon > 0$ and δ the function $\epsilon k - \delta k^\lambda$ is monotonically increasing from the first point where it becomes positive. This fact may easily be verified by differentiation. The base case follows.

For the inductive step we let $k > k_0$ and assume that the statement holds for all values at most k. To prove the statement for k, let us consider a graph $G \in \mathcal{G}$ and a treewidth-η-modulator S of G of size at most k. If $|S| < k$ then by the induction hypothesis, G has a $((1+\epsilon)(k-1) - \delta(k-1)^\lambda, r)$-protrusion decomposition. Since the function $(1+\epsilon)k - \delta k^\lambda$ is non-decreasing whenever it is non-negative and $k - 1 \geq k_0$ it follows that this is also a $((1+\epsilon)k - \delta k^\lambda, r)$-protrusion decomposition. Thus we may assume that $|S| = k$.

By Observation 3.5, the treewidth of G is at most $\mathrm{tw}(G) \leq \alpha \cdot k^\lambda$. We apply Lemma 2.1 to (G, S) and obtain a 2/3-balanced separation (A_1, A_2) of (G, S) of order at most $\mathrm{tw}(G) + 1 \leq \alpha \cdot k^\lambda + 1$. Let $L = A_1 \setminus A_2$, $X = A_1 \cap A_2$ and $R = A_2 \setminus A_1$. Since (A_1, A_2) is a 2/3-balanced separation, it follows that there exists a real $\frac{1}{3} \leq a \leq \frac{2}{3}$ such that $|L \cap S| \leq a|S|$ and $|R \cap S| \leq (1-a)|S|$.

Consider now the graph $G[L \cup X]$. Let $S_L = (S \cap L) \cup X$. Clearly S_L is a treewidth-η-modulator of $G[L \cup X]$ and $|S_L| \leq a|S| + |X| \leq ak + 2\alpha k^\lambda$. Since $k > k_0$, the properties of k_0 imply that $\frac{2}{3}k + 2\alpha k^\lambda \leq k - 1$. Furthermore,

$$|S_L| \geq |S \setminus R| \geq a|S| \geq \frac{1}{3}k \geq \frac{1}{3}k_0 \geq 1.$$

Thus we may apply the induction hypothesis to $G[L \cup X]$, with treewidth modulator S_L and obtain a $((1+\epsilon)|S_L| - \delta|S_L|^\lambda, r)$-protrusion decomposition $\mathcal{P}_L$ of $G[L \cup X]$ with core containing S_L.

We now consider the graph $G[R \cup X]$ and define $S_R = (S \cap R) \cup X$. We have that

$$(1-a)|S| \leq |S_R| \leq (1-a)|S| + |X| \leq (1-a)k + 2\alpha k^\lambda.$$

Similarly, we have a $((1+\epsilon)|S_R| - \delta|S_R|^\lambda, r)$-protrusion decomposition $\mathcal{P}_R$ of $G[R \cup X]$ with core containing S_R.

From the protrusion decompositions $\mathcal{P}_L$ and $\mathcal{P}_R$ we construct a protrusion decomposition $\mathcal{P}$ of G. The core of $\mathcal{P}$ is the union of the cores of $\mathcal{P}_L$ and $\mathcal{P}_R$. The set of protrusions of $\mathcal{P}$ is the union of the set of protrusions of $\mathcal{P}_L$ and $\mathcal{P}_R$ respectively. Since the cores of $\mathcal{P}_L$ and $\mathcal{P}_R$ contain X the protrusions of $\mathcal{P}_L$ and $\mathcal{P}_R$ are also protrusions in G. Therefore $\mathcal{P}$ is a $((1+\epsilon)(|S_L|+|S_R|) - \delta(|S_L|^\lambda + |S_R|^\lambda), r)$-protrusion decomposition of G containing S in its core. Thus, to finish the proof it is sufficient to show that

$$(1+\epsilon)(|S_L| + |S_R|) - \delta(|S_L|^\lambda + |S_R|^\lambda) \leq (1+\epsilon)k - \delta k^\lambda.$$

Since $|S_L| \leq ak + 2\alpha k^\lambda$ and $|S_R| \leq (1-a)k + 2\alpha k^\lambda$ it follows that $|S_L| + |S_R| \leq k + 4\alpha k^\lambda$ and that $\delta(|S_L|^\lambda + |S_R|^\lambda) \geq \delta k^\lambda(a^\lambda + (1-a)^\lambda)$. It follows that

$$\begin{aligned}
&(1+\epsilon)(|S_L| + |S_R|) - \delta(|S_L|^\lambda + |S_R|^\lambda) \\
&\leq (1+\epsilon)(k + 4\alpha k^\lambda) - \delta k^\lambda(a^\lambda + (1-a)^\lambda) \\
&= (1+\epsilon)k - \delta k^\lambda + k^\lambda(4\alpha(1+\epsilon) - \delta(a^\lambda + (1-a)^\lambda - 1)) \\
&\leq (1+\epsilon)k - \delta k^\lambda + k^\lambda(4\alpha(1+\epsilon) - \rho\delta) \\
&= (1+\epsilon)k - \delta k^\lambda
\end{aligned}$$

In the transition from the third to the fourth line we used that $(1-a)^\lambda + a^\lambda - 1 \geq \rho$ for any a between $\frac{1}{3}$ and $\frac{2}{3}$. This completes the induction.

We have now proved that for any $k \geq \frac{1}{3}k_0$, any graph $G \in \mathcal{G}$ such that $G \in \mathcal{G}$ has a treewidth-η-modulator S of size at most k, then G has $((1+\epsilon)k - \delta k^\lambda, r)$-protrusion decomposition, which is also a $((1+\epsilon)k, r)$-protrusion decomposition. If G has a treewidth-η-modulator $S \neq \emptyset$ and $|S| \leq \frac{k_0}{3}$ then the protrusion decomposition with core S and one protrusion, $V(G) \setminus S$ is a $(|S|, \eta)$-protrusion decomposition. Since $|S| \leq (1+\epsilon)|S|$ and $\eta \leq r$ this completes the proof of the lemma. □

Lemmata 3.6, 3.7, and 3.10 imply the following theorem, which is the first main technical contribution of this paper.

Theorem 3.11. *Let $\mathcal{G}$ be a graph class with the SQGM property and Π be a minor-bidimensional linear-separable problem. There exists a constant c such that, if G and $k \in \mathbb{Z}^+$ are the graph and the parameter of a yes-instance of $\Pi \Cap \mathcal{G}$, then G admits a $(c \cdot k, c)$-protrusion decomposition.*

Similarly, if $\mathcal{G}$ is a graph class with the SQGC property and Π be a contraction-bidimensional linear-separable problem, then for every $(G, k) \in \Pi \Cap \mathcal{G}$, G admits a $(c \cdot k, c)$-protrusion decomposition.

Proof. Let (G,k) be a yes-instance of a minor-bidimensional linear-separable problem Π with $G \in \mathcal{G}$. By Lemma 3.6, for any real $\epsilon > 0$ there exists an integer $\eta \geq 0$ such that graph G has a treewidth-η-modulator S of size at most $\epsilon \cdot k$. Then by Lemma 3.10, G has $(c \cdot k, c)$-protrusion decomposition for some choice of constant c.

The proof for contraction-bidimensional problems is similar, the only difference that we use Lemma 3.7 instead of Lemma 3.6. □

4 Finite index and finite integer index

In this section we prove the second main technical contribution of the paper: "CMSO + separability $\Rightarrow$ finite integer index".

4.1 Definitions on boundaried graphs

Boundaried graphs. A *boundaried graph* $\mathbf{G} = (G, B, \lambda)$ is a triple consisting of a graph G, a set $B \subseteq V(G)$ of distinguished vertices and an injective labelling λ from B to the set $\mathbb{Z}_{>0}$. The set B is called the *boundary* of $\mathbf{G}$ and the vertices in B are called *boundary vertices* or *terminals*. Also, the graph G is the *underlying graph* of $\mathbf{G}$. Given a boundaried graph $\mathbf{G} = (G, B, \lambda)$, we define its *label set* by $\Lambda(\mathbf{G}) = \{\lambda(v) \mid v \in B\}$. If $\mathbf{G} = (G, B, \lambda)$ is a boundaried graph and $S \subseteq V(G)$, then the pair $\hat{\mathbf{G}} = (\mathbf{G}, S)$ is an *annotated boundaried graph* and we say that G is the *underlying graph* of $\hat{\mathbf{G}}$, B is its *boundary* (labeled by λ), and S is its *annotated set*.

Given a finite set $I \subseteq \mathbb{Z}_{>0}$, we define $\mathcal{F}_I$ (resp. $\hat{\mathcal{F}}_I$) as the class of all boundaried graphs (resp. annotated boundaried graphs) whose label set is I. Similarly, we define $\mathcal{F}_{\subseteq I} = \bigcup_{I' \subseteq I} \mathcal{F}_{I'}$ (resp. $\hat{\mathcal{F}}_{\subseteq I} = \bigcup_{I' \subseteq I} \hat{\mathcal{F}}_{I'}$). We also denote by $\mathcal{F}$ (resp. $\hat{\mathcal{F}}$) the class of all boundaried graphs, i.e., $\mathcal{F} = \mathcal{F}_{\subseteq \mathbb{Z}_{>0}}$ (resp. $\hat{\mathcal{F}} = \hat{\mathcal{F}}_{\subseteq \mathbb{Z}_{>0}}$). Finally we say that a boundaried graph $\mathbf{G}$ is a *t-boundaried* graph if $\Lambda(\mathbf{G}) \subseteq \{1, \ldots, t\}$.

Let $\mathcal{G}$ be a class of (not boundaried) graphs. We say that a boundaried graph $\mathbf{G}$ *belongs to* $\mathcal{G}$ if the underlying graph of $\mathbf{G}$ belongs to $\mathcal{G}$. We also use $V(\mathbf{G})$ to denote the vertex set of the underlying graph of $\mathbf{G}$.

The gluing operation. Let $\mathbf{G}_1 = (G_1, B_1, \lambda_1)$ and $\mathbf{G}_2 = (G_2, B_2, \lambda_2)$ be two boundaried graphs. We denote by $\mathbf{G}_1 \oplus \mathbf{G}_2$ the graph (not boundaried) obtained by taking the disjoint union of G_1 and G_2 and identifying equally-labeled vertices of the boundaries of G_1 and G_2. The gluing operation maintains edges of both graphs that are glued, i.e., in $\mathbf{G}_1 \oplus \mathbf{G}_2$ there is an edge between two labeled vertices if there is either an edge between them in G_1 or in G_2, (in case of multi-graphs, multiplicities of edges are summed up in the new graph).

Let $G = \mathbf{G}_1 \oplus \mathbf{G}_2$ where $\mathbf{G}_1 = (G_1, B_1, \lambda_1)$ and $\mathbf{G}_2 = (G_2, B_2, \lambda_2)$ are boundaried graphs. We define the *glued* set of $\mathbf{G}_i$ as the set $B_i^\star = \lambda_i^{-1}(\Lambda(\mathbf{G}_1) \cap \Lambda(\mathbf{G}_2)), i = 1, 2$. For a vertex $v \in V(\mathbf{G}_1)$ we define its heir $\mathsf{heir}(v)$ in G as follows: if $v \notin B_1^\star$ then $\mathsf{heir}(v) = v$, otherwise $\mathsf{heir}(v)$ is the result of the identification of v with an equally labeled vertex in $\mathbf{G}_2$. The *heir* of a vertex in G_2 is defined symmetrically. The common boundary of $\mathbf{G}_1$ and $\mathbf{G}_2$ in G is equal to $\mathsf{heir}(B_1^\star) = \mathsf{heir}(B_2^\star)$ where the evaluation of h on vertex sets is defined in the obvious way.

Let now $\hat{\mathbf{G}}_1 = (\mathbf{G}_1, S_1)$ and $\hat{\mathbf{G}}_2 = (\mathbf{G}_2, S_2)$ be two annotated boundaried graphs. We define $\hat{\mathbf{G}}_1 \oplus \hat{\mathbf{G}}_2$ as the annotated graph $\hat{G} = (\mathbf{G}_1 \oplus \mathbf{G}_2, \mathsf{heir}(S_1) \cup \mathsf{heir}(S_2))$.

4.2 Finite Index

Let ϕ be a predicate on annotated graphs. In other words ϕ is a function that takes as input a graph G and vertex set $S \subseteq V(G)$, and outputs true or false. We define a *canonical equivalence relation* $\equiv_\phi$ on boundaried annotated graphs as follows. For two annotated boundaried graphs $\hat{\mathbf{G}}_1 = (\mathbf{G}_1, S_1)$ and $\hat{\mathbf{G}}_2 = (\mathbf{G}_2, S_2)$, we say that $\hat{\mathbf{G}}_1 \equiv_\phi \hat{\mathbf{G}}_2$ if $\Lambda(\mathbf{G}_1) = \Lambda(\mathbf{G}_2)$ and for every annotated boundaried graph $\hat{\mathbf{G}} = (\mathbf{G}, S)$ we have that

$$\phi(\hat{\mathbf{G}}_1 \oplus \hat{\mathbf{G}}) = \mathsf{true} \Leftrightarrow \phi(\hat{\mathbf{G}}_2 \oplus \hat{\mathbf{G}}) = \mathsf{true}$$

It is easy to verify that $\equiv_\phi$ is an equivalence relation. We say that ϕ is *finite state* if, for every finite $I \subseteq \mathbb{Z}_{>0}$, the equivalence relation $\equiv_\phi$ has a finite number of equivalence classes when restricted to $\hat{\mathcal{F}}_I$. A formal proof of the following fact can be found in [6, 13].

Proposition 4.1. *For every CMSO-definable predicate ϕ on annotated graphs, ϕ has finite state.*

Given a CMSO-definable predicate ϕ on annotated graphs and an $I \subseteq \mathbb{Z}_{>0}$, we use the notation $\hat{\mathcal{R}}_{\phi,I}$ for a set containing one minimum-size representative from each of the equivalence classes of $\equiv_\phi$ when restricted to $\hat{\mathcal{F}}_I$.

4.3 Finite Integer Index

Definition 4.2. [**Canonical equivalence on boundaried graphs.**] Let Π be a parameterized graph problem whose instances are pairs of the form (G, k). Given two boundaried graphs $\mathbf{G}_1, \mathbf{G}_2 \in \mathcal{F}$, we say that $\mathbf{G}_1 \equiv_\Pi \mathbf{G}_2$ if $\Lambda(\mathbf{G}_1) = \Lambda(\mathbf{G}_2)$ and there exists a *transposition constant* $c \in \mathbb{Z}$ such that

$$\forall (\mathbf{F}, k) \in \mathcal{F} \times \mathbb{Z} \qquad (\mathbf{G}_1 \oplus \mathbf{F}, k) \in \Pi \Leftrightarrow (\mathbf{G}_2 \oplus \mathbf{F}, k + c) \in \Pi.$$

Note that the relation $\equiv_\Pi$ is an equivalence relation. Observe that c could be negative in the above definition. This is the reason we gave the definition of parameterized problems to include negative parameters also.

Notice that two boundaried graphs with different label sets belong to different equivalence classes of $\equiv_\Pi$. We are now in position to give the following definition.

Definition 4.3. [**Finite Integer Index**] A parameterized graph problem Π whose instances are pairs of the form (G, k) has *Finite Integer Index* (or simply has *FII*), if and only if for every finite $I \subseteq \mathbb{Z}^+$, the number of equivalence classes of $\equiv_\Pi$ that are subsets of $\mathcal{F}_I$ is finite.

4.4 A condition for proving FII

Theorem 4.4. *If Π is a separable CMSO-minimization or CMSO-maximization problem, then Π has FII.*

Proof. We prove the statement for vertex subset CMSO-minimization problems. The proofs for CMSO-maximization problems and the edge subset variants are identical. Let ψ be the CMSO-predicate defining Π, that is $(G,k) \in \Pi$ if and only if there exists a vertex set S of size at most k such that $(G,S) \models \psi$. Let I be a finite subset of $\mathbb{Z}_{>0}$ and let $\hat{\mathcal{R}}_{\psi,I}$ be a set of representatives of $\equiv_\psi$ when restricted to $\hat{\mathcal{F}}_I$, the set of boundaried graphs with label set I. Let f be a function such that Π is f-separable.

Given a boundaried graph $\mathbf{G} \in \mathcal{F}_I$, we define the function $\xi_{\mathbf{G}} : \hat{\mathcal{R}}_{\psi,I} \to \mathbb{Z}_{\geq 0} \cup \{\perp\}$ such that if $(\mathbf{G}^*, S^*) \in \hat{\mathcal{R}}_{\psi,I}$, then

$$\xi_{\mathbf{G}}(\mathbf{G}^*, S^*) \quad = \quad \min\{|S| \mid (\mathbf{G}, S) \equiv_\psi (\mathbf{G}^*, S^*)\}. \tag{1}$$

Here $\xi_{\mathbf{G}}(\mathbf{G}^*, S^*) = \perp$ if no set S such that $(\mathbf{G}, S) \equiv_\psi (\mathbf{G}^*, S^*)$ exists. We may think of $\xi_{\mathbf{G}}$ as a partial function where $\perp$ means that the function is left undefined. We also define function $\chi_{\mathbf{G}}(\mathbf{G}^*, S^*) : \hat{\mathcal{R}}_{\psi,I} \to \mathbb{Z} \cup \{\perp\}$ as follows

$$\chi_{\mathbf{G}}(\mathbf{G}^*, S^*) = \begin{cases} \xi_{\mathbf{G}}(\mathbf{G}^*, S^*) - OPT_\Pi(G), & \text{if } \xi_{\mathbf{G}}(\mathbf{G}^*, S^*) \in [OPT_\Pi(G) - f(|I|), OPT_\Pi(G) + f(|I|)] \\ \perp, & \text{otherwise.} \end{cases}$$

Thus $\chi_{\mathbf{G}}$ outputs $\perp$ if $\xi_{\mathbf{G}} = \perp$ or when $\xi_{\mathbf{G}}(\mathbf{G}^*, S^*) \notin [OPT_\Pi(G) - f(|I|), OPT_\Pi(G) + f(|I|)]$. Clearly $\chi_{\mathbf{G}}(\mathbf{G}^*, S^*) \in [-f(|I|), f(|I|)] \cup \{\perp\}$ for all choices of $(\mathbf{G}^*, S^*) \in \hat{\mathcal{R}}_{\psi,I}$. This means that $\chi_{\mathbf{G}}$ may take at most $2 \cdot f(|I|) + 2$ different values, and hence there are at most $(2 \cdot f(|I|) + 2)^{|\hat{\mathcal{R}}_{\psi,I}|}$ different functions $\chi_{\mathbf{G}}$.

Now, we define an equivalence relation $\sim$ on $\mathcal{F}_I$ such that, given $\mathbf{G}_1, \mathbf{G}_2 \in \mathcal{F}_I$,

$$\mathbf{G}_1 \sim \mathbf{G}_2 \text{ if and only if } \chi_{\mathbf{G}_1} = \chi_{\mathbf{G}_2}$$

and keep in mind that $\sim$ has a finite number of equivalent classes. We will prove that $\sim$ is a refinement of $\equiv_\Pi$. For this we prove that if $\mathbf{G}_1 \sim \mathbf{G}_2$, then $\mathbf{G}_1 \equiv_\Pi \mathbf{G}_2$.

We assume that $\mathbf{G}_1 \sim \mathbf{G}_2$, where $\mathbf{G}_1 = (G_1, B_1)$ and $\mathbf{G}_2 = (G_2, B_2)$. We also set $c = OPT_\Pi(G_2) - OPT_\Pi(G_1)$ (notice that c depends only on $\mathbf{G}_1$ and $\mathbf{G}_2$). In what follows, we prove that, for every $(\mathbf{F}, k) \in \mathcal{F}_I \times \mathbb{Z}$, $(\mathbf{G}_1 \oplus \mathbf{F}, k) \in \Pi$ implies that $(\mathbf{G}_2 \oplus \mathbf{F}, k + c) \in \Pi$ (the direction $(\mathbf{G}_1 \oplus \mathbf{F}, k) \notin \Pi \Rightarrow (\mathbf{G}_2 \oplus \mathbf{F}, k + c) \notin \Pi$ is symmetric).

We denote $H = \mathbf{G}_1 \oplus \mathbf{F}$ and we consider some $(\mathbf{F}, k) \in \mathcal{F}_I \times \mathbb{Z}$. Let S be an optimal solution of H. That is, $|S| = OPT_\Pi(H)$ and $(H, S) \models \psi$. The fact that $(H, k) \in \Pi$ means that the size of S is at most k. We define $S_1 = S \cap V(\mathbf{G}_1)$ and $S^\star = S \setminus S_1$ and we know that $(\mathbf{G}_1 \oplus \mathbf{F}, S_1 \cup S^\star) \models \psi$. Notice that $|S_1| + |S^\star| = |S_1 \cup S^\star| \leq k$.

Claim: There exists no $S_1' \subseteq V(\mathbf{G}_1)$ such that $|S_1'| < |S_1|$ and $(\mathbf{G}_1, S_1') \equiv_\psi (\mathbf{G}_1, S_1)$.
Proof of Claim: Suppose in contrary that such an S_1' exists. Since $(\mathbf{G}_1, S_1') \equiv_\psi (\mathbf{G}_1, S_1)$, we have that $(\mathbf{G}_1 \oplus \mathbf{F}, S_1' \cup S^\star) \models \psi$. However, this implies that $|S_1' \cup S^\star| = |S_1'| + |S^\star| < |S_1| + |S^\star| = OPT_\Pi(H)$, which is a contradiction. This concludes the claim. □

Let $(\mathbf{G}', S') \in \hat{\mathcal{R}}_{\psi,I}$ such that $(\mathbf{G}', S') \equiv_\psi (\mathbf{G}_1, S_1)$. By the above claim and (1), we have that $\xi_{\mathbf{G}_1}(\mathbf{G}', S') = |S_1|$. Let $A_1 = \mathsf{heir}(V(\mathbf{G}_1))$ and $A_2 = \mathsf{heir}(V(\mathbf{F}))$ (here by the heir of a vertex set we mean the union of all vertex heir it contains) . Here, function heir is defined with respect to the operation $H = \mathbf{G}_1 \oplus \mathbf{F}$ and certainly (A_1, A_2) is a separation of H of order $|I|$.

Recall that $S_1 = S \cap A_1$ and $|S| = OPT_\Pi(H)$. The f-separability of Π implies that

$$|S \cap A_1| \in [OPT_\Pi(H[A_1]) - f(|I|), OPT_\Pi(H[A_1]) + f(|I|)],$$

and therefore,

$$\xi_{\mathbf{G}_1}(\mathbf{G}', S') \in [OPT_\Pi(G_1) - f(|I|), OPT_\Pi(G_1) + f(|I|)].$$

Then $\chi_{\mathbf{G}_1}(\mathbf{G}', S') \neq \perp$ and thus $\chi_{\mathbf{G}_1}(\mathbf{G}', S') = |S_1| - OPT_\Pi(G_1)$. As $\mathbf{G}_1 \sim \mathbf{G}_2$, this means that $\chi_{\mathbf{G}_2}(\mathbf{G}', S') = |S_1| - OPT_\Pi(\mathbf{G}_1)$. Since $\chi_{\mathbf{G}_2}(\mathbf{G}', S') \neq \perp$ it follows that $\chi_{\mathbf{G}_2}(\mathbf{G}', S') = \xi_{\mathbf{G}_2}(\mathbf{G}', S') - OPT_\Pi(G_2)$. We conclude that

$$\xi_{\mathbf{G}_2}(\mathbf{G}', S') = |S_1| + c. \tag{2}$$

By the definition of ξ and (2), there exists a set S_2 of size at most $|S_1| + c$ such that $(\mathbf{G}_2, S_2) \equiv_\psi (\mathbf{G}', S')$. As $(\mathbf{G}_1, S_1) \equiv_\psi (\mathbf{G}_2, S_2)$, the fact that $(\mathbf{G}_1 \oplus \mathbf{F}, S_1 \cup S^\star) \models \psi$, implies that $(\mathbf{G}_2 \oplus \mathbf{F}, S_2 \cup S^\star) \models \psi$. But then $|S_2 \cup S^\star| = |S_2| + |S^\star| \leq |S_1| + c + |S^\star| \leq k + c$. This means that $(\mathbf{G}_2 \oplus \mathbf{F}, k + c) \in \Pi$. Hence $\mathbf{G}_1 \equiv_\Pi \mathbf{G}_2$.

Thus for every pair of boundaried graphs $\mathbf{G}_1$ and $\mathbf{G}_2$, condition $\mathbf{G}_1 \sim \mathbf{G}_2$, yields that $\mathbf{G}_1 \equiv_\Pi \mathbf{G}_2$. Since $\sim$ has a finite number of equivalent classes, we conclude that Π has FII. □

5 Proof of the main theorem: Putting things together

Now everything prepared to pipeline the results about protrusion decomposition and FII with the framework from [6]. We start from the following definitions.

Definition 5.1. [(f, a)**-protrusion replacement family**] Let Π be a parameterized graph problem, let $f : \mathbb{Z}^+ \to \mathbb{Z}^+$ be a non-decreasing function and let $a \in \mathbb{Z}^+$. An (f, a)-*protrusion replacement family* for Π is a collection $\mathcal{A} = \{\mathsf{A}_i \mid i \geq 0\}$ of algorithms, such that algorithm A_i receives as input a pair (I, X), where

- I is an instance of Π whose graph and parameter are G and $k \in \mathbb{Z}$,
- X is an i-protrusion of G with at least $f(i) \cdot k^a$ vertices,

and outputs an equivalent instance I^* such that, if G^* and k^* are the graph and the parameter of I^*, then $|V(G^*)| < |V(G)|$ and $k^* \leq k$.

The following two properties for a parameterized graph problem Π were defined in [6].

A [**Protrusion replacement**:] There exists an (f, a)-protrusion replacement family $\mathcal{A}$ for Π, for some function $f : \mathbb{Z}^+ \to \mathbb{Z}^+$ and some $a \in \mathbb{Z}^+$.

B [**Protrusion decomposition:**] There exists a constant c such that, if G and $k \in \mathbb{Z}^+$ are the graph and the parameter of a yes-instance of Π then G admits a $(c \cdot k, c)$-protrusion decomposition.

Our kernelization result is based on the following master theorem from [6].

Theorem 5.2. *If a parameterized graph problem* Π *has property* **A** *for some nonnegative constant* a *and property* **B** *for some constant* c*, then* Π *admits a kernel of size* $O(k^{a+1})$.

The following lemma proven in [6], shows that every parameterized problems with FII admits $(f, 0)$-protrusion replacement families.

Lemma 5.3. *Every parameterized graph problem* Π *that has FII has the protrusion replacement property* **A** *for* $a = 0$.

Proof of Theorem 1.1. Let Π be a CMSO-definable linear-separable minor-bidimensional problem on an H-minor-free graph class $\mathcal{G}$ for some fixed graph H. By Proposition 2.10, $\mathcal{G}$ has SQGM property.

By Theorem 4.4, every separable CMSO-minimization or CMSO-maximization problem Π has FII. Thus by Lemma 5.3, Π has the protrusion replacement property **A** for $a = 0$. By Theorem 3.11, Π has the protrusion decomposition property **B**. By Theorem 5.2, Π admits a linear kernel.

The proof for contraction-bidimensional problems is almost identical. □

6 Conclusion

We conclude with the discussions on the relation between the new and previous kernelization meta-theorems, the running times of kernelization algorithms and an open question.

Bidimensionality vs. quasi-coverability. Let us finally note how the techniques developed to prove the kernelization meta-theorem in this paper, can be used to refine the results from [6].

Let r be a non-negative integer. We say that a parameterized graph problem Π has the *radial r-coverability property* if all yes-instances of Π encode graphs embeddable in some surface of Euler genus at most r and there exist such an embedding of G and a set $S \subseteq V(G)$ such that $|S| \leq r \cdot k$ and $\mathbf{R}^r_G(S) = V(G)$. A parameterized graph problem Π has the *radial r-quasi-coverability property* if all yes-instance s of Π encode graphs embeddable in some surface of Euler genus at most r and there exist such an embedding and a set $S \subseteq V(G)$ such that $|S| \leq r \cdot k$ and $\mathrm{tw}(G - \mathbf{R}^r_G(S)) \leq r$. Every problem Π that has the radial r-quasi-coverability property is radially r-quasi-covervable. The converse is not necessarily true.

The main two meta-theorems from [6] say that quasi-coverable problems with finite integer index admit linear kernels on graphs of bounded genus and that coverable CMSO-definable problems admit polynomial kernels on graphs of bounded genus. In order to refine [6], we need the following lemma.

Lemma 6.1. *If* Π *is a quasi-coverable problem on graphs of bounded genus, then there exists a bidimensional separable problem* Π^* *such that if* $(G, k) \in \Pi$*, then* $OPT_{\Pi^*}(G) = O(k)$.

The lemma follows from the following observation. Let Π be an r-quasi-coverable problem. Consider the following problem Π^*: for a graph G of genus g, pair $(G, k) \in \Pi^*$ if and only if there is a subset of vertices $S \subseteq V(G)$ such that $|S| \leq r \cdot k$ and $\mathrm{tw}(G \setminus \mathbf{R}^r_G(S)) \leq r$. Because Π is r-quasi-coverable, we have that if $(G, k) \in \Pi$, then $OPT_{\Pi^*}(G) \leq r \cdot k$. As in the case with Treewidth-η-Modulator, it is easy to see that Π^* is contraction-bidimensional and separable.

By Lemma 3.10, there is a protrusion decomposition for Π^* and thus there is a decomposition for Π. This implies one of the main result of [6]: Every quasi-coverable problem with finite integer index admits a linear kernel on graphs of bounded genus. This also can be used to show that every quasi-coverable CMSO-definable problem admits a polynomial kernel on graphs of bounded genus. Finally, let us remark that the results for quasi-coverable problems on graphs of bounded genus from [6] can be extended to graphs excluding some apex graph as a minor. Here, one needs to adapt the notion of covering a graph. Instead of covering graph by balls in the metric space of its radial-graph, we can use the metric space of some supergraph H of a given graph G, which excludes as a minor the same apex graph as graph G.

Running time. In Theorem 1.1 we do not specify the running time of the kernelization algorithm. The running time of such algorithms depends on the running time hidden in the master theorem from [6] (Theorem 5.2), which in turn, depends on how fast one can identify and replace protrusions in a graph. By applying recent "protrusion replacer" from [28], we obtain that the running time of kernelization algorithm in Theorem 1.1 is linear in the input size.

Open question. We conclude with the following open question. For separable contraction-bidimensional CMSO-definable problems, the technique developed in this paper yields the existence of a linear kernel on H-minor-free graphs only when H is an apex graph. An interesting open question is to identify general logic and combinatorial conditions for contraction-bidimensional problems which yield a polynomial kernelization on H-minor-free graphs. For some contraction-bidimensional problems, like Dominating Set or Connected Dominating Set, linear kernels are known to exist for H-minor-free graphs [31], H-topological-minor-free graphs [32], and even for graphs of bounded expansion [25] but we are still far from a general understanding of the complexity of kernelization for such problems beyond apex-minor-free graphs.

References

[1] J. Alber, M. R. Fellows, and R. Niedermeier. Polynomial-time data reduction for dominating set. *J. ACM*, 51(3):363–384, 2004.

[2] J. Alber, H. Fernau, and R. Niedermeier. Parameterized complexity: exponential speed-up for planar graph problems. *J. Algorithms*, 52(1):26–56, 2004.

[3] S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *J. Algorithms*, 12(2):308–340, 1991.

[4] H. L. Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theoretical Computers Science*, 209(1-2):1–45, 1998.

[5] H. L. Bodlaender, F. V. Fomin, D. Lokshtanov, E. Penninkx, S. Saurabh, and D. M. Thilikos. (Meta) Kernelization. In *Proceedings of the 50th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 629–638. IEEE, 2009.

[6] H. L. Bodlaender, F. V. Fomin, D. Lokshtanov, E. Penninkx, S. Saurabh, and D. M. Thilikos. (Meta) Kernelization. *CoRR*, abs/0904.0727, 2013. To appear in J. of ACM.

[7] H. L. Bodlaender and E. Penninkx. A linear kernel for planar feedback vertex set. In *Proceedings of the 3rd International Workshop on Parameterized and Exact Computation (IWPEC)*, volume 5018 of *Lecture Notes in Comput. Sci.*, pages 160–171. Springer, 2008.

[8] H. L. Bodlaender, E. Penninkx, and R. B. Tan. A linear kernel for the k-disjoint cycle problem on planar graphs. In *Proceedings of the 19th International Symposium on Algorithms and Computation (ISAAC)*, volume 5369 of *Lecture Notes in Comput. Sci.*, pages 306–317. Springer, 2008.

[9] C. Chekuri and J. Chuzhoy. Polynomial bounds for the Grid-Minor Theorem. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 60–69. ACM, 2014.

[10] J. Chen, H. Fernau, I. A. Kanj, and G. Xia. Parametric duality and kernelization: Lower bounds and upper bounds on kernel size. *SIAM J. Computing*, 37(4):1077–1106, 2007.

[11] J. Chen, I. A. Kanj, and W. Jia. Vertex cover: further observations and further improvements. *J. Algorithms*, 41(2):280–301, 2001.

[12] B. Courcelle. The monadic second-order logic of graphs I: Recognizable sets of finite graphs. *Information and Computation*, 85:12–75, 1990.

[13] B. Courcelle. The expression of graph properties and graph transformations in monadic second-order logic. *Handbook of Graph Grammars*, pages 313–400, 1997.

[14] M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.

[15] M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.

[16] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Bidimensional parameters and local treewidth. *SIAM J. Discrete Math.*, 18(3):501–511, 2004.

[17] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Fixed-parameter algorithms for (k, r)-center in planar graphs and map graphs. *ACM Transactions on Algorithms*, 1(1):33–47, 2005.

[18] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs. *J. ACM*, 52(6):866–893, 2005.

[19] E. D. Demaine and M. Hajiaghayi. Bidimensionality: new connections between FPT algorithms and PTASs. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 590–601. SIAM, 2005.

[20] E. D. Demaine and M. Hajiaghayi. Graphs excluding a fixed minor have grids as large as treewidth, with combinatorial and algorithmic applications through bidimensionality. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 682–689. ACM-SIAM, 2005.

[21] E. D. Demaine and M. Hajiaghayi. The bidimensionality theory and its algorithmic applications. *Comput. J.*, 51(3):292–302, 2008.

[22] E. D. Demaine and M. Hajiaghayi. Linearity of grid minors in treewidth with applications through bidimensionality. *Combinatorica*, 28(1):19–36, 2008.

[23] F. Dorn, F. V. Fomin, and D. M. Thilikos. Subexponential parameterized algorithms. *Computer Science Review*, 2(1):29–39, 2008.

[24] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity.* Texts in Computer Science. Springer, 2013.

[25] P. G. Drange, M. Dregi, F. V. Fomin, S. Kreutzer, D. Lokshtanov, M. Pilipczuk, M. Pilipczuk, F. Reidl, F. S. Villaamil, S. Saurabh, S. Siebertz, and S. Sikdar. Kernelization and Sparseness: the Case of Dominating Set. In *Proceedings of the 33rd Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 47 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 31:1–31:14, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[26] J. Flum and M. Grohe. *Parameterized Complexity Theory.* Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin, 2006.

[27] F. V. Fomin, P. A. Golovach, and D. M. Thilikos. Contraction obstructions for treewidth. *J. Combinatorial Theory Ser. B*, 101(5):302–314, 2011.

[28] F. V. Fomin, D. Lokshtanov, N. Misra, M. Ramanujan, and S. Saurabh. Solving d-SAT via backdoors to small treewidth. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (S)DA)*, pages 630–641. SIAM, 2015.

[29] F. V. Fomin, D. Lokshtanov, V. Raman, and S. Saurabh. Bidimensionality and EPTAS. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 748–759. SIAM, 2011.

[30] F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Bidimensionality and kernels. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 503–510. SIAM, 2010.

[31] F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Linear kernels for (connected) dominating set on H-minor-free graphs. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 82–93. SIAM, 2012.

[32] F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Linear kernels for (Connected) Dominating Set on graphs with excluded topological subgraphs. In *Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 20 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 92–103. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013.

[33] J. Guo and R. Niedermeier. Invitation to data reduction and problem kernelization. *SIGACT News*, 38(1):31–45, 2007.

[34] J. Guo, R. Niedermeier, and S. Wernicke. Fixed-parameter tractability results for full-degree spanning tree and its dual. In *Proceedings of the 2nd International Workshop on Parameterized and Exact Computation (IWPEC)*, volume 4169 of *Lecture Notes in Comput. Sci.*, pages 203–214. Springer, 2006.

[35] I. A. Kanj, M. J. Pelsmajer, G. Xia, and M. Schaefer. On the induced matching problem. In *Proceedings of the 25th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 08001, pages 397–408. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, Berlin, 2008.

[36] S. Kratsch. Recent developments in kernelization: A survey. *Bulletin of EATCS*, 2(113), 2014.

[37] D. Lokshtanov, M. Mnich, and Saurabh. Linear kernel for planar connected dominating set. In *Proceedings of Theory and Applications of Models of Computation, (TAMC 2009)*, Lecture Notes in Comput. Sci. Springer, 2009.

[38] N. Misra, V. Raman, and S. Saurabh. Lower bounds on kernelization. *Discrete Optim.*, 8(1):110–128, 2011.

[39] H. Moser and S. Sikdar. The parameterized complexity of the induced matching problem in planar graphs. In *Proceedings First Annual International WorkshopFrontiers in Algorithmics (FAW)*, volume 4613 of *Lecture Notes in Comput. Sci.*, pages 325–336. Springer, 2007.

[40] R. Niedermeier. *Invitation to fixed-parameter algorithms*, volume 31 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, Oxford, 2006.

[41] N. Robertson and P. D. Seymour. Graph minors. III. Planar tree-width. *J. Combinatorial Theory Ser. B*, 36:49–64, 1984.

[42] N. Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *J. Combinatorial Theory Ser. B*, 41(1):92–114, 1986.

[43] S. Thomassé. A $4k^2$ kernel for feedback vertex set. *ACM Transactions on Algorithms*, 6(2), 2010.